

PRACTICAL ORBIT DETERMINATION FOR AEROBRAKING WITH ACCUMULATED ACCELEROMETER DATA

Brian Young*

Navigation during aerobraking missions has historically relied solely on radiometric Doppler data. Accelerometer data, available on all modern spacecraft, has the potential to supplement Doppler data, increasing operational tempos, filling tracking data gaps, and improving the accuracy of orbit determination solutions. However, integration of Doppler and accelerometer data into the batch filter solutions used for operational navigation has proven to be practically challenging, due to the narrow region of adequate linear approximation for acceleration incurred by drag. This work proposes to use the accumulated velocity changes derived from these values as a more compatible data type. Demonstrations of results focus on NASA’s Maven and ESA’s TGO mission, and how this technique can be used to improve operations for the upcoming Maven aerobraking phase with reduced tracking data and lighter staffing requirements.

INTRODUCTION

Aerobraking is an efficient method to reduce the apoapsis altitude of a spacecraft’s orbit, using atmospheric drag at periapsis to accumulate a significant velocity reduction over hundreds of drag passes. The first operational use of aerobraking was by the Magellan mission to Venus in May 1993¹. NASA’s previous three Mars orbiters before Maven, Mars Global Surveyor (1997)², Mars Odyssey (2001)³, and the Mars Reconnaissance Orbiter (2006)⁴ also used aerobraking to reduce propellant requirements and achieve their final near-circular science orbits. The European Space Agency (ESA) performed experimental aerobraking with their Venus Express mission in 2014⁵, and used the method to achieve the science orbit for the ExoMars Trace Gas Orbiter (TGO) in 2017–2018^{6,7}. NASA’s Maven mission, which used propulsive maneuvers to achieve its final science orbit in 2014, passes through thinner but still drag-inducing altitudes in its multi-year operational science orbit, with occasional week-long “deep dips” to deeper altitudes for additional scientific studies⁸, leading to many of the same challenges faced by aerobraking missions⁹. Finally, Maven has been requested by NASA to reduce its apoapsis altitude, via aerobraking at deep dip altitudes, to better provide relay support for NASA’s assets at Mars’ surface, starting in March 2019. Supporting this aerobraking task with limited tracking data and staffing levels is the primary impetus for this work.

During aerobraking, and in similar regimes, the Navigation team is charged with three tasks: reconstructing the trajectory, predicting the upcoming trajectory, and selecting

*Navigation Engineer, Mission Design and Navigation Section, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109

apoapsis maneuvers to alter the periapsis altitude and maintain acceptable atmospheric density levels. Because drag is the most significant non-gravitational perturbation on the trajectory, estimating and predicting that drag is the dominant consideration in all three tasks. Operationally, a reference density model, either an exponential model or a more complex model such as MarsGRAM¹⁰, is used with a per-orbit multiplicative density scale factor to capture unmodelled or random variations in the actual density levels. Reconstructions produce a history of scale factors, which are then averaged to compute a predicted scale factor, which in turn informs maneuver selection. On Mars, it is typically assumed that the scale factor can vary by over 100% between orbits, and that the predicted value may be in error by as much as 30% (3σ); the Venusian atmosphere is more stable but still unpredictable. This navigation/maneuver process is usually performed on a daily cycle for aerobraking missions and Maven deep dips, and weekly cycle for Maven during normal operations. During Maven aerobraking, the planned cycle will likely last 2–3 days, as a compromise between workload and aerobraking efficiency.

Radiometric Doppler, provided by NASA’s Deep Space Network (DSN) or ESA’s ESTRACK, is the primary data type used for orbit determination (OD) on missions orbiting rocky planets. When present, these data indicate the period and timing of each orbit, so that the drag ΔV , which is proportional to the scale factor, can be reconstructed accurately. While this has been sufficient for past missions, it is limiting in two ways. First, for aerobraking missions, operational timelines have little margin, so faster reconstruction is helpful, but it takes some time to accumulate enough Doppler data after a drag pass to accurately reconstruct the corresponding scale factor. Second, accurate reconstruction of drag ΔV requires nearly continuous coverage by the ground stations, since at most only two periapses can occur between tracking passes before the ability to distinguish the effects of individual periapsis passes is lost, though the average of the scale factor across that gap can still be computed. For Maven, where science observations during these tracking gaps require reconstruction uncertainties under 3 km, the loss of accuracy during these gaps is also a challenge, especially during periods with high drag levels. While DSN/ESTRACK coverage is a limited resource, historically aerobraking missions have been able to request continuous coverages. However, because Maven has been in this class of orbit for an extended period of time, coverage is limited to approximately one 8 hour pass per day, except during deep dips. For the upcoming Maven aerobraking, a significant challenge is that tracking will not be continuous either.

Most modern spacecraft have inertial measurement units (IMUs) on-board that include accelerometers to measure high-resolution data of non-gravitational forces applied to the vehicle, such as drag or thruster firings. This provides a direct measurement of the perturbing force, complementing Doppler data, and potentially mitigating the two given challenges. Because the data would be available as soon as they could be downlinked, they could be rapidly integrated into an OD solution, improving timeline margins for daily operations. The acceleration data can also be stored onboard the spacecraft until the next available downlink opportunity, and can therefore be used to fill tracking data gaps and allow more accurate trajectory reconstructions. For these reasons, methods to integrate accelerometer data with the OD process have been pursued for years. For reasons discussed below, however, practical use has proven stubbornly difficult, so much so that Maven continues to not use the accelerometer data for OD three years after the start of science operations.

Still, accelerometer data have seen a great deal of other practical use in aerobraking-style missions. Periapsis timing estimators, which update on-board ephemerides and command times based on accelerometer-derived periapsis times, were used for Mars Reconnaissance Orbiter (MRO) aerobraking, and are a key enabling technology for Maven science operations. Detailed reconstructions from Maven accelerometer data have provided new insights into atmospheric structures^{11,12}. Even within the navigation task, the data have proven useful, even if they cannot be integrated into the filtering process. Comparisons of Doppler-based reconstructions to the measured accelerations are used to validate that solutions are correct before delivery to other users. Comparing the observed accelerometer data to the expected drag ΔV can be used to better determine a priori estimates of the scale factors, improving filter performance. ESA teams have also used the data as a force model directly, as an alternative to providing it as a measured value.

Still, integration of the measurements into operational OD would be the most straightforward approach. Focusing on Maven, this integration would improve reconstruction accuracy and navigation team's ability to observe trends in atmospheric behavior, and could significantly improve the team's ability to operate the upcoming aerobraking portion of the mission. This work will demonstrate an approach to reduce the available data to its most important components in a way that works well with existing processes, providing all of the advantages described here.

IMPEDIMENTS TO HIGH RATE ACCELEROMETER MEASUREMENT USAGE

Deep space navigation at JPL relies on batch processing of measurements with a linear least squares filter. Currently implemented by the Monte software¹³, this process has remained fundamental to navigation since the earliest deep space missions, with the legacy Orbit Determination Program (ODP) operating on the same principles. A set of initial conditions and parameterized force models are used to generate a trajectory, and then, for the available set of measurements, the expected set of observables corresponding to the a priori parameters are computed. Given partial derivatives of the computed measurement values to all the estimable parameters, the set of parameter updates that minimize the residual difference between the observed and computed values, in a linear approximation, are calculated using a least squares minimization to generate updated model parameters. This process is repeated with the new parameter values until the updates converge to near zero. This batch approach is preferred to a sequential estimation process like a Kalman filter because it allows the entire reconstructed trajectory to be computed at once, and ensures that the final solution can be reproduced without explicit reference to the measurements.

For aerobraking-class missions, the primary estimable parameters are the initial condition and a per-orbit scale factor. Additional terms include solar pressure, reaction wheel desaturation impulsive ΔV s, and extra periapsis impulses to account for lift, sideslip, and peak offsets in the drag impulse. Typically, the initial condition is well known based on preceding reconstructions or short single-orbit fits, and the additional terms are relatively small. Thus, the primary uncertainty is in the scale factors, or the expected anti-velocity ΔV at each periapsis due to drag, which manifest linearly as variations in the mean anomaly and orbital period.

Previous work by Jones et. al. for Maven has pursued use of acceleration data as direct measurements for use in the filter, dividing the high-rate ΔV data by the accumulation

time to get accelerations, and filtering those acceleration values against continuous non-gravitational forces in the model (primarily drag) to improve estimates during the orbit¹⁴. Since the accelerometer data are proportional to the drag force, direct measurements of that acceleration should easily correct the scale factor. However, operational use has proven extremely difficult for a few reasons. First, the noise on a magnitude measurement follows a non-central chi distribution, breaking filter assumptions of normal random variables. Second, biases in the data are especially difficult to remove with the noise terms, which themselves appear as a bias. Third, the actual drag around periapsis often does not follow the shape of the curves in the modeled atmosphere, so that no good fit with a single scale factor per periapsis is possible. Solutions for these challenges are available, including exclusion of data below the known noise floor, pre-processing of vector data before conversion to magnitudes, use of vector data rather than magnitude data, and/or appropriate weighting of the measurements.

However, the most significant problem facing the use of high-rate accelerometer data is that linearization breaks down with small timing errors in the initial estimates. As a comparison, consider Doppler, a measure of the Earth-line velocity, which varies sinusoidally with true anomaly. Fitting Doppler data estimates the period and periapsis times of each orbit, matching the phase and wavelength of the Doppler signature. Because these phase and periapsis times are related to the drag ΔV via linear relationships, corrections to timing and period are predictably implemented as scale factor changes, so that the filter converges to a solution with reliable scale factors for the covered periapses. Given that the characteristic variation of the Doppler signature varies over an orbital period, intuition indicates that fitting with Doppler data can handle significant timing errors in the initial estimate, perhaps up to tens of minutes for even the shortest two-hour orbits. Other model parameters shrink this region of linearity, but practical experience shows that minutes of error can be accommodated by this process.

Instead, high-rate accelerometer data directly measure the drag acceleration, rather than the timing shifts caused by the accumulated ΔV , and residual differences between the observed and modeled measurements can be caused either by an incorrect drag level at the current periapsis, or by incorrect timing of the drag impulse caused by errors in the previous periapses. When the timing errors are small, this is handled well, because an error in the current scale factor will appear as residual errors in the same direction, while a timing error will show an increasing or decreasing trend in the residuals, so that the appropriate correction can be determined from the ensemble. However, if the modeled peak drag is further from the measured values, then the timing error will instead appear as a large scale factor error, since the low drag near the modeled tails will need to be increased drastically to reach the measured values; this large and invalid correction will then further disrupt the estimates of later periapses, leading to a failure to converge. Given that the drag impulses of interest to this study have a full-width at half maximum of two to three minutes, one would expect that timing errors larger than approximately one minute would encounter these problems. Practical experience shows this to be true, with timing errors over approximately 30 seconds proving difficult to correct. This time range is an order of magnitude smaller than what Doppler can correct, and shorter than the variations seen in operational cases, leading to the significant challenges seen using these data.

There are techniques to deal with this. First fitting the Doppler data, and only then

adding accelerometer data could mitigate the issue significantly, as long as data gaps were short enough to maintain reasonable timing uncertainty in those gaps. Incrementally fitting one or two periapses at a time can also be used to mitigate the problem, since all previous periapses will have been well-fit, reducing any potential timing errors for the next periapsis in the sequence. However, this incremental fitting process may still struggle if telemetry downlink problems prevent some earlier periapses from being received properly. Heuristic processes to guess appropriate a priori values could also help. However, these options are time consuming, add additional steps, break with long-standing practices, and increase the risk associated with telemetry data not being downlinked, so that the benefits have been judged less helpful than the costs. They also require human experience, experimentation, and evaluation to use effectively, reducing the ability of analysts to complete regular operation tasks quickly and reliably, or to improve workflows through automation.

Jah, et. al., demonstrated an alternative technique, using the navigation batch filter to generate a solution up to the start of the drag pass, and then passing that state to a sequential filter to include accelerometer data through periapsis¹⁵. This process was then repeated for subsequent orbits, building up OD solutions over time as telemetry data was received, with a goal of quickly generating solutions using only accelerometer data, before sufficient Doppler data was collected. As in the incremental fitting process, this is successful because it disentangles the ambiguity of timing and scale factor corrections, keeping the timing errors at the time of accelerometer processing minimized. Like the other proposed techniques, this breaks processes associated with the batch model and is sensitive to missed telemetry during Doppler gaps, preventing it from being considered for operational usage on subsequent missions.

Ultimately, while Doppler data allows the navigation batch filter to reliably converge with a broad range of initial conditions, high-resolution acceleration measurements are only reliable in a narrow range due the sensitivity of their partial derivatives to the timing of the orbit. While there are many theoretical methods to handling this problem, none of these fit within the framework used for normal operations, and thus have not seen significant use.

PROPOSED ACCUMULATED ΔV MEASUREMENT

Given the diagnosis that high-rate accelerometer data are too sensitive to timing errors, and knowing that the typical data resolution gives far more detail of the force profile than can be fit using typical engineering models of the atmosphere, an alternative approach is suggested. The accelerometer data is of interest for navigation because it indicates the total ΔV incurred during a given drag pass. Therefore, accumulating the high-rate ΔV s into a single velocity shift over a given time period will provide the desired information in a reduced format less sensitive to a priori timing uncertainty. For Maven and TGO, a time period of 10 minutes before and after the peak drag is used, because this encompasses the drag pass with sufficient margin for expected timing uncertainty, without including any wheel desaturation impulses or other events that might occur away from periapsis.

This accumulation could take the form of either a vector summation or a total magnitude. The vector data will provide information on any lift or sideslip incurred during the drag pass, but for high-altitude passes where the signal is weak relative to noise and biases, the off-velocity components may be too small to be reliable. While deweighting the data to appropriate levels would ensure they would not over-constrain the filtered solution, on

principle it is preferable to use a simpler measurement with only the useful information. Additionally, though the primary interest is the total ΔV , the effective time at which that ΔV is applied may also be valuable. This time, the effective “centroid”, akin to a “center of mass” for a given impulse, can provide information on longitudinal density gradients or other terms that shift the time of the maximum acceleration, and in extreme cases may provide constraints on periapsis times to further aid convergence*.

To generate the observables, consider the set of raw accelerometer data as a N -long sequence of vectors \mathbf{v}_i that are each centered at a time t_i , measured as seconds from an arbitrary epoch such as the measured peak or a priori periapsis. These vectors can be in any reference frame, usually either an inertial frame or the spacecraft body-fixed reference frame. Given these, the accumulated vector value for a set of accelerometer data is

$$\mathbf{v} = \sum \mathbf{v}_i, \quad (1)$$

and the magnitude is

$$v = \sqrt{\mathbf{v} \cdot \mathbf{v}}. \quad (2)$$

Note that one should not use the sum of the magnitudes of \mathbf{v}_i , since for a normally distributed vector, the magnitude is chi-distributed, with a non-zero mean that can accumulate to large errors in the derived observable. Finally, the centroid of the drag pulse is computed as

$$t = \frac{\sum t_i (\mathbf{v}_i \cdot \hat{\mathbf{v}})}{\sum \mathbf{v}_i \cdot \hat{\mathbf{v}}}. \quad (3)$$

Note that the high-rate vectors are projected onto the accumulated unit vector $\hat{\mathbf{v}} = \mathbf{v}/v$ to create a scalar measurement that is normally-distributed, minimizing the effects of noise for this term as well.

There are five primary sources of error in the raw accelerometer data: biases, scale factors, non-orthogonality, spacecraft rotations, and noise. Biases are straightforward to remove, because the time before periapsis is well-known to have negligible non-gravitational acceleration, and thus any level of readout can be assumed to be a bias. This is often removed on-board the spacecraft and does not need to concern the analyst. However, for Maven science operations, because the signal is much weaker than for other missions and phases, it is necessary to remove the bias on the ground based on both pre- and post-periapsis quiet periods before computing the summations. The scale factor cannot be detected solely from accelerometer data, and removal relies on estimation by the filter, where the Doppler provides independent measurements of the drag ΔV between two time periods. Non-orthogonality is not considered in this work, but its existence should be acknowledged as a potential unmodeled source of error. Rotational motion can usually be removed on-board the spacecraft, given knowledge of attitude rates and the accelerometer location, but the analyst should ensure they understand the behavior of their system, and understand that mismodeled centers of mass can create small signatures. Finally, noise manifests as uncertainty in the final observed value, and this uncertainty should be used to inform the weights applied to the measurements.

*Alternatively, for these cases, a better approach might be to provide this time directly as a measurement of periapsis or peak drag time

To quantify the expected noise in the accumulated measurements, assume that each vector component of \mathbf{v}_i is normally distributed and independent, with standard deviation σ . The vector sum is thus the sum of these independent variables, so each component is normally distributed with standard deviation

$$\sigma_v = \sqrt{N}\sigma. \quad (4)$$

The magnitude is the root sum square of these three normally-distributed vector components, and follows a non-central chi-squared distribution. However, assuming that the total summation is significantly larger than the expected noise (i.e. $v \gg \sigma_v$), this is approximately Gaussian, with minimal bias and standard deviation σ_v . If this condition does not hold, the magnitude value should be treated with caution. The distribution of the centroid can be computed as

$$\begin{aligned} t &= \frac{\sum t_i \mathcal{N}(v_i, \sigma^2)}{\sum \mathcal{N}(v_i, \sigma^2)} \\ &= \frac{\mathcal{N}(\sum t_i v_i, \sigma^2 \sum t_i)}{\mathcal{N}(\sum v_i, N\sigma^2)} \\ &\approx \mathcal{N}\left(\frac{\sum t_i v_i}{\sum v_i}, \frac{\sum t_i^2}{\sum v_i} \sigma^2\right), \end{aligned}$$

where $\mathcal{N}(\mu, \sigma^2)$ represents a normally distributed variable, and with the final step assuming that $v \gg \sigma_v$. Thus, assuming sufficient signal quality, the centroid has a standard deviation of

$$\sigma_t = \sqrt{\frac{\sum t_i^2}{\sum v_i}} \sigma. \quad (5)$$

Finally, it should be noted that the normal and independent assumption given here may not be valid. Because of correlated noise values, stochastic biases, or unmodeled rotational motion, values may need to be dewighted more based on experience and analyst judgment.

Plots of the acceleration values, accumulated ΔV , and centroid times are shown in figure 1 for Maven science operations, near 0.1 kg/km^3 , in figure 2 for a Maven Deep Dip orbit, near 3 kg/km^3 , and in figure 3 for TGO aerobraking, near 20 kg/km^3 . These charts show the acceleration values, the accumulated ΔV , and the value of $\sum t_i v_i / v$ as they accumulate over time, for all three body-fixed axes and for the total magnitude. Comparing these, first note the relative noise level for Maven science operations that complicates usage. Noise levels indicate that the off-axis terms are poorly known, although the total ΔV is reasonably certain. Additionally, the centroid time is exceptionally noisy, because the weighted average is especially sensitive to the noisy tails of the drag impulse. However, the stronger signals in the other two cases reduce these problems drastically, showing useful and believable data for all three vector terms as well as the centroid. Another item of interest is that while it is possible to compute per-axis centroid times, even with moderately strong signals these can be noisy and unreliable, justifying the use of a single value computed along the total ΔV vector. Also noticeable in the two stronger signals is that the fine structure of the acceleration is quite variable compared to an ideal Gaussian curve; use of this technique ignores this fine structure, which is not usually estimable using navigation models. Finally, the TGO data show significant impulses shortly after periapsis, which are autonomous

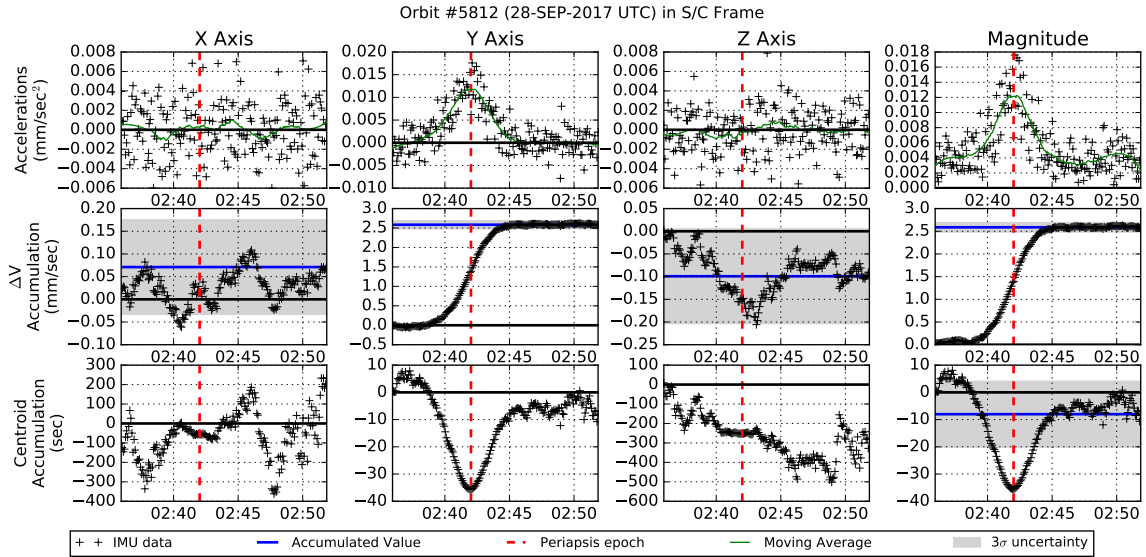


Figure 1. Sample accelerometer data for Mars Reconnaissance Orbiter nominal science

thruster firings for attitude control. Because they are also modeled in the trajectory they will be properly accounted for in the final filter solution, though shortening the accumulation time to exclude these impulses is also possible.

For both Mars Reconnaissance Orbiter and Mars Global Surveyor, finding the data associated with a periapsis is straightforward, since both missions use on-board periapsis timing estimators and only enable high-rate collection during appropriate times, so that clusters of high-rate points are easily identified as valid drag impulses. Other missions or other acceleration sources may present different challenges, so peak-finding algorithms may be necessary to find the appropriate data.

The computation of model-derived measurement values should follow the same formulae, with additional terms to represent biases and scale factors. For simplicity, assume continuous forces and impulsive ΔV s can be discretized into a sequence of high-rate modeled vectors $\bar{\mathbf{v}}_i$ that apply at times \bar{t}_i . This formulation assumes constant values for the bias vector β and the scalar or per-axis scale factor α for the duration of a measurement, though time-dependent variations are possible and are a simple extension. With these definitions, the modeled vector sum is

$$\bar{\mathbf{v}} = \alpha \sum \bar{\mathbf{v}}_i + \beta (t_f - t_0). \quad (6)$$

The modeled magnitude measurement can be computed as

$$\bar{v} = \sqrt{\bar{\mathbf{v}} \cdot \bar{\mathbf{v}}}. \quad (7)$$

Finally, the centroid time is

$$\bar{t} = \frac{\alpha \sum \bar{t}_i \bar{\mathbf{v}}_i \cdot \hat{\bar{\mathbf{v}}} + \frac{1}{2} (\beta \cdot \hat{\bar{\mathbf{v}}}) (t_f - t_0)^2}{\alpha \sum \bar{\mathbf{v}}_i \cdot \hat{\bar{\mathbf{v}}} + (\beta \cdot \hat{\bar{\mathbf{v}}}) (t_f - t_0)}. \quad (8)$$

These model-derived values should be computed using the same coordinate frame definitions and reference epoch values as the observed data.

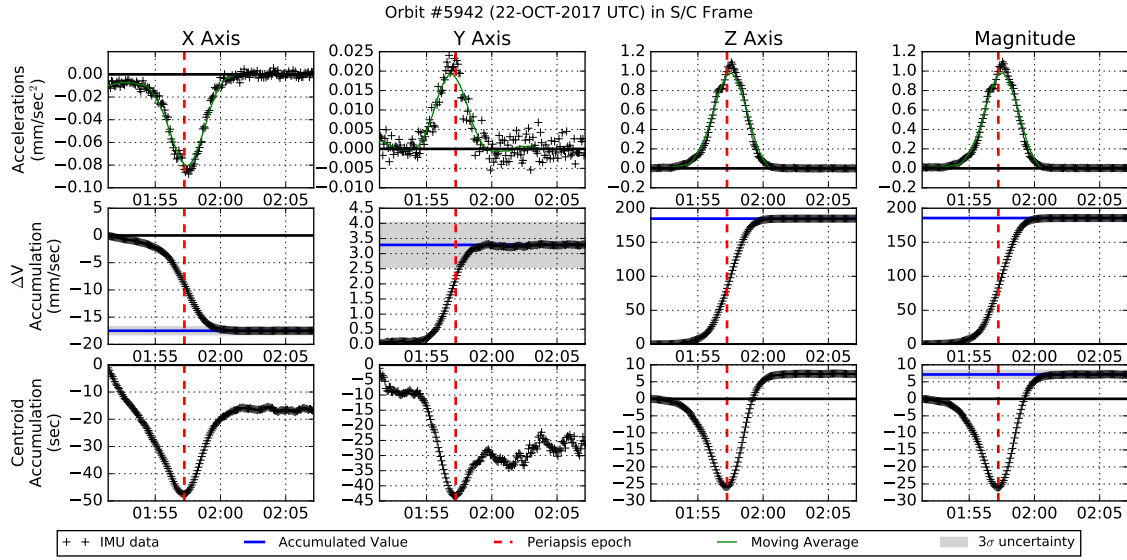


Figure 2. Sample accelerometer data for Maven deep dips

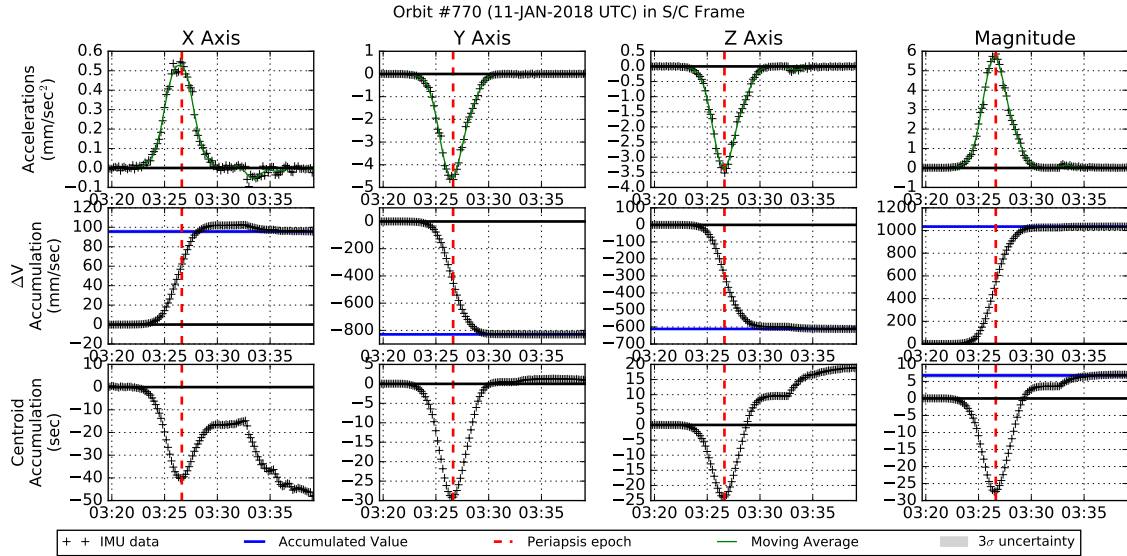


Figure 3. Sample accelerometer data for TGO aerobraking

APPLICATION TO OPERATIONAL MISSIONS

These derived measurement types have been integrated with the operational navigation system using hooks for arbitrary user-defined implementations, based on the formulae above. This allows the standard tools used for operational OD to be used with this new data type, facilitating comparisons between cases and evaluations of practical use cases. The implementation considers acceleration due to drag and impulses used to model thruster firings and off-axis drag models, and can include bias and scale factor terms. All of these cases estimate the initial state with a priori values from previous reconstructs, a per-orbit scale factor on a MarsGRAM atmosphere model with the Map Year 0 settings, desaturation impulses at telemetry-derived locations, solar pressure with a single scale factor for the entire arc, impulses at periapsis in the cross-track and radial directions to model lift and sideslip due to drag, and negatively-correlated velocity-direction impulses 5 minutes before and after periapsis to capture centroid shifts due to density gradients. Accelerometer scale factors are estimated per-axis and per-periapsis using a loose 5% 1σ uncertainty. Biases are not estimated since they can be removed more effectively by pre-processing the data.

Maven Nominal Science

The first case to consider is from Maven nominal science, in mid-March 2018, covering seven periapses (6725 through 6731) with tracking data up to the first periapsis, and surrounding the final periapsis, so that the tracking gap includes six periapses. Three versions of this arc were analyzed, a Doppler-only baseline, a case with the magnitude accelerometer data, and a case with the full vector accelerometer data, all weighted at 0.2 mm/sec, since tighter weights induced unrealistic results and poor convergence. Centroid data were not considered due to the large noise terms indicated in the previous section. All three versions of this case converged in three iterations of the filter, with corrected parameter values within expected ranges. Considering the results, of particular interest for maneuver planning is the reconstruction of density scale factors to aid in the prediction of future densities. Figure 4 shows the estimated scale factors (with 3σ uncertainty bounds) for the three cases. The Doppler-only case in this chart shows a typical pattern, where the estimated values fall along a linear trend, matching the phase and period before and after the tracking gap, with large uncertainties, and it is understood that while the average of these values is accurate, the specific values are under-determined. As expected, the addition of accelerometer magnitude data yields superior per-periapsis estimates, but the vector data provide little new information.

An important related metric is the reconstructed trajectory position uncertainty, plotted in figure 5. Improvements to the accuracy of the trajectory may allow improved scientific results, since the instruments will have more information about where they were actually pointing during observations. While the Doppler-only baseline is within agreed-upon requirements for accuracy, improvements are always welcomed. These results show that the downtrack uncertainty, which is dominated by drag, is reduced by an order of magnitude, from 350 meters near periapsis to less than 50 meters, near the 10 meter level achieved during Doppler passes in this regime. Similar improvements occur in the radial direction, which is primarily a function of improved knowledge of the semimajor axis. The additional vector data provide little benefit in the downtrack and radial components. Interestingly, the vector data is shown to increase overall cross-track uncertainties from 70 meters to




Figure 4. Estimated density scale factors for Maven science operations

110 meters; the reason for this is not fully understood, but may be an artifact of particular modeling choices. Overall, this demonstrates that the simple magnitude data provide significant benefit, with only modest improvement with the additional vector data.

The long term trend and variability in the scale factor is also of interest. While the average value during data gaps is sufficient to compute a predicted average, the lack of specific values in coverage gaps has two consequences. First, the measured variability is artificially reduced, and this variability can affect decision making on where to maneuver within the density corridor. Second, there may be repeating patterns within the atmosphere that indicate waves in the atmosphere that could not be easily detected without reasonable per-orbit scale factors. A batch of 400 orbits, spanning periapses 6400 to 6800 were recomputed using these measurements, with the original and updated scale factor estimates shown in figure 6. Additionally, this exercise demonstrated that the measurements were effective across a range of potential orbits, and could be used effectively by automated processing techniques.

Maven Deep Dip

Next, consider a case from a Maven deep dip, where the altitude was reduced to 120 km to perform in situ studies of densities near 3 kg/km^3 . The case spans orbits 5931–5937 in late October 2017. During this time, there were no significant tracking data gaps, except those at periapsis and apoapsis, due to occultations and spacecraft attitude, so that scale factors could be reliably reconstructed. For this case, the addition of accelerometer data provided modest gains in accuracy, reducing the 1σ downtrack uncertainty from 50 meters to 40 meters, as shown in figure 7, with even smaller gains from the addition of centroid data. Note that in this case, the cross-track uncertainty exhibits expected behavior, with vector data improving cross-track accuracy. More interesting is that typically, this long

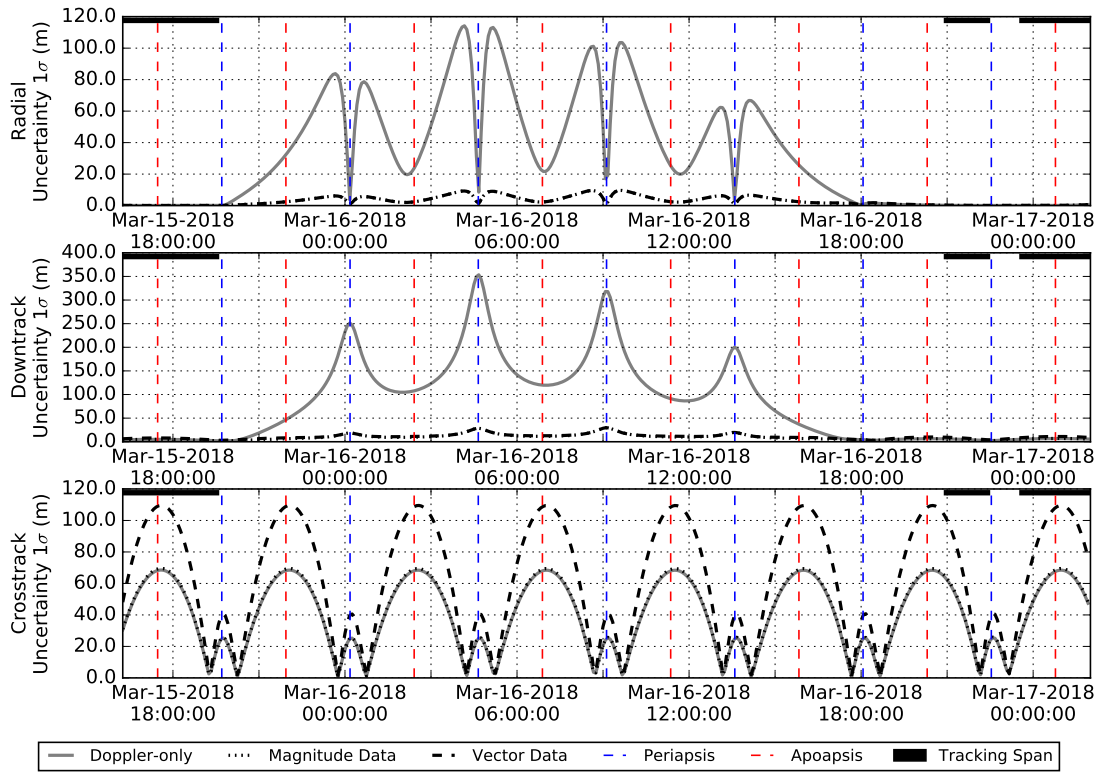


Figure 5. Position post-fit uncertainties for Maven science operations




Figure 6. Updated density scale factor estimates for Maven orbits 6400–6800

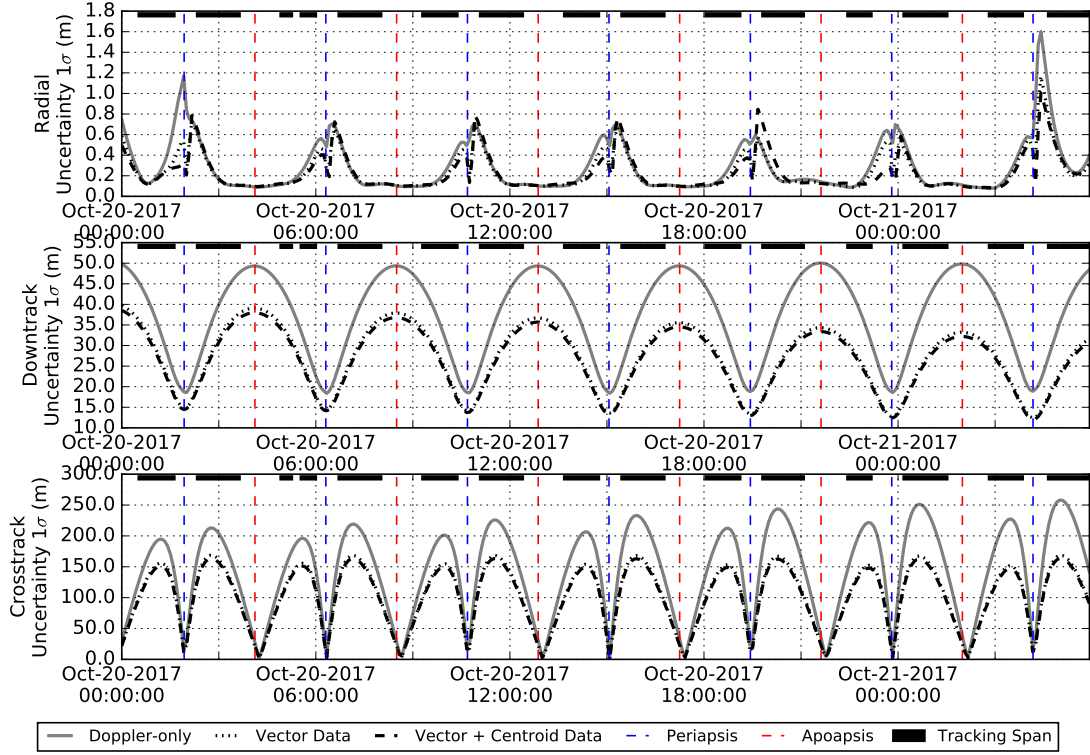


Figure 7. Position post-fit uncertainties for MAVEN deep dip

of an arc with this level of drag perturbation would not converge without incrementally fitting individually periapses. However, the addition of the accelerometer data allowed the solution to converge from a poor initial state, albeit haphazardly, showing these data to have stabilizing influence on filter behavior.

An alternative scenario where data covering the middle five orbits is removed is also considered, as an analogue for expected conditions during planned aerobraking. The timing drift over this long of an arc means that the filter will not converge without special measures, and the process of incrementally fitting periapses will fail since there are no data covering five periapses. Instead a fit of the two Doppler-covered periapses, and an average scale factor guess across the gap computed through the known total period shift would be necessary without any available accelerometer data. The addition of accelerometer-derived ΔV measurements makes the problem tractable, allowing an incremental fit to operate as expected. With the accelerometer data included, results are similar to the nominal science cases, with position uncertainties reduced from 10 km to under 1 km, and radial errors reduced from 3.5 km to under 250 m, as shown in figure 8, with the associated improvements in scale factor estimates shown in figure 9. Trajectory differences with the full-Doppler case are shown in figure 10, demonstrating that the accelerometer data reasonably match the Doppler data. Note that the addition of the centroid data has a significant effect on the cross-track uncertainty; additionally, better matching of centroid times will allow better measurement of atmospheric waves, which manifest as shifts in the peak drag epoch,

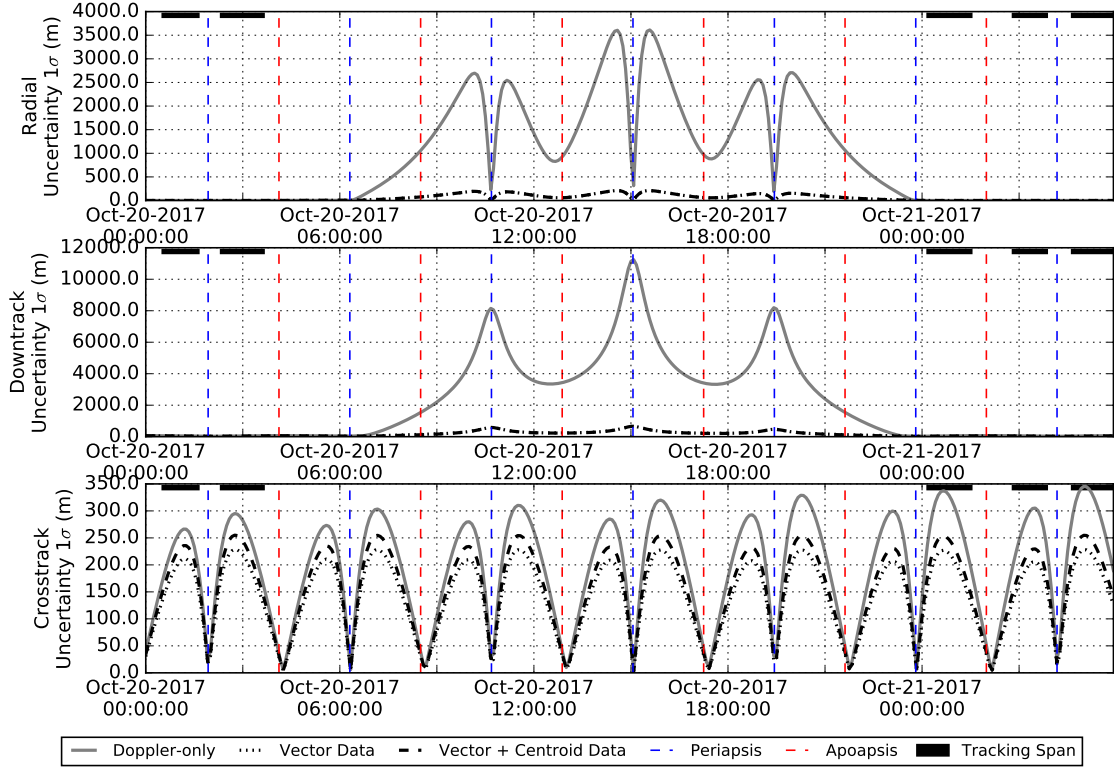


Figure 8. Position post-fit uncertainties for Maven aerobraking scenario

recommending the use of these data in deep dip/aerobraking operations.

TGO Aerobraking

ESA’s ExoMars TGO mission used aerobraking to reduce the initial 24 hour orbit to a 2 hour orbit from March 2017 through February 2018, with a short recess for solar conjunction in August 2017. During the later stages, as the period dropped below 6 hours, the European Space Operations Center (ESOC) team utilized this method to incorporate accelerometer data as a ΔV vector, while the JPL team, performing shadow navigation, had the capability to do so but did not regularly choose to do so, due to continuous Doppler coverage and archival challenges associated with user-defined measurement types. However, results were comparable between the two teams throughout the mission, within the loose range considered acceptable during this science observation-free period. In February 2018 as the orbits shortened, a tracking pass covering multiple orbits was missed, making Doppler fits spanning the gap difficult. The addition of accelerometer data performed exactly as expected in this case, helping define each of these unobserved periapses, allowing a viable reconstruction to be generated, and operations to continue unimpeded.




Figure 9. Estimated density scale factors for Maven aerobraking scenario




Figure 10. Trajectory differences between aerobraking scenario and true full-Doppler case

CONCLUSIONS AND FUTURE WORK

Use of onboard IMU-derived accelerometer data within the operational orbit determination process in aerobraking-class missions has been pursued for years, with team members struggling to develop a process that fit within established procedures and worked reliably. The key problem was the poor linearization of short term accelerations with respect to multi-orbit arcs, suggesting the use of longer-term accumulated ΔV s, along with their effective centroid epoch, to capture the dominant effects of the drag impulse while maintaining compatibility with existing procedures. Tests in a variety of operational cases in the Maven nominal science orbit and during deep dips show that these accumulated ΔV values can be integrated without degrading the analyst’s ability to generate a timely solution, and often improve convergence in addition to the primary goal of improving reconstructions during tracking data gaps by an order of magnitude. Furthermore, use of these methods during TGO aerobraking by both ESOC and JPL teams demonstrate their effectiveness.

Looking forward to the aerobraking that Maven will begin in March 2019, new challenges emerge. This endeavor will involve flying to deep dip altitudes, but with sparser tracking schedules, less frequent deliveries, and reduced staffing, over a period of two to three months. Use of these data greatly improves the Navigation team’s ability to perform this task successfully. First, the ability to fill tracking gaps will make filtering over multiple untracked periapses a less tedious and haphazard process, so that solutions will be more reliable and less time consuming, and more suitable for automation. Second, this improved knowledge of the trajectory within the data gaps will allow significantly more science observations to be performed than might otherwise be possible with the large uncertainties associated with Doppler-only solutions. Third, better knowledge of particular drag passes allows a better understanding of trends that may affect decision making. Ultimately, while the aerobraking task could be completed with current capabilities, the ability to include IMU-derived data reduces the risks associated with difficult-to-fit and time-consuming solutions during an already tense operations process.

Future work will consider using this technique to observe other forces. Reaction wheel desaturation thrustings occur regularly, 10–15 minutes after periapsis for Maven, and are poorly characterized in telemetry downlinked by the spacecraft. Being small and near periapsis, the effects of these thrustings are also difficult to distinguish from other dynamic effects. Accelerometer data during these events are noisy, which has dissuaded previous consideration of inclusion. However, use of this technique could allow new higher quality reconstructions of these events, improving the fit and the characterization of the spacecraft system. Similarly, this could be done with executed maneuvers, with a focus on improving convergence, since the events are already well-characterized by Doppler. Applications to other classes of missions may exist as well.

Finally, for Maven aerobraking, a further inquiry should pursue use of downlinked periapsis timing estimator data. This data, usually including ΔV and timing data, is used to track periapsis times in on-board command sequences. Because the accumulated values are computed on-board, the data volume is reduced and is downlinked on the low-gain antenna, which is used more commonly than the high-gain antenna since it does not disrupt scientific observations. If these low-rate data prove of sufficient quality, they could serve two purposes. First, they would be more timely than the full set of telemetry values, allowing their use for the most recent orbits, including in automated “quicklook” runs. Second, during

aerobraking, it remains to be seen if Maven can schedule MSPA (multiple spacecraft per aperture) DSN passes while a station is communicating with other missions. These passes only allow downlink, and thus cannot provide two-way Doppler, but could still provide this low-rate telemetry, enabling the constant stream of data available for missions with continuous tracking coverage.

Ultimately, this technique provides a powerful and straightforward way to include onboard IMU data with orbit determination, and can aid future aerobraking tasks, including the upcoming aerobraking by the Maven mission.

ACKNOWLEDGMENTS

This work was carried out at the Jet Propulsion Laboratory, California Institute Technology, under a contract with the National Aeronautics and Space Administration.

Copyright 2018 California Institute of Technology. Government sponsorship acknowledged.

REFERENCES

- [1] Daniel T. Lyons, R. Stephen Saunders, and Douglas G. Griffith. The magellan venus mapping mission: Aerobraking operations. *Acta Astronautica*, 35(9):669 – 676, 1995. Challenges of Space for a Better World.
- [2] P. Esposito, V. Alwar, P. Burkhart, S. Demcak, E. Graat, M. Johnston, and B. Portock. Navigating mars global surveyor through the martian atmosphere: Aerobraking 2, paper presented at aas/aiaa astrodynamic specialist conference, 1999.
- [3] Robert A. Mase, Peter Antreasian, Julia L. Bell, Tomas L. Martin-Mur, and John C. Smith Jr. Mars odyssey navigation experience. *Journal of Spacecraft and Rockets*, 42(3):386–393, 2018/04/16 2005.
- [4] Stacia Long, Tung-Han You, C. Halsell, Ramachand Bhat, Stuart Demcak, Eric Graat, Earl Higa, Dolan Highsmith, Neil Mottinger, and Moriba Jah. Mars Reconnaissance Orbiter Aerobraking Navigation Operation. American Institute of Aeronautics and Astronautics, 2018/02/26 2008.
- [5] Saturnino VAL SERRA, Olivier BONNAMY, Olivier WITASSE, and Octavio CAMINO. Venus express aerobraking. *IFAC Proceedings Volumes*, 44(1):715 – 720, 2011. 18th IFAC World Congress.
- [6] F. Castellini, B. Godard, and G. Bellei. Flight Dynamics Operational Experience from Exomars TGO aerobraking campaign at Mars. In *15th Spaceops*, 2018.
- [7] M. Denis, P. Schmitz, R. Sangiorgi, P. Mitschdoerfer, M. Montagna, H. Renault, and N. Kutrowski. Thousand times through the atmosphere of mars: Aerobraking the exomars trace gas orbiter. In *15th Spaceops*, 2018.
- [8] B. M. Jakosky, R. P. Lin, J. M. Grebowsky, J. G. Luhmann, D. F. Mitchell, G. Beutelschies, T. Priser, M. Acuna, L. Andersson, D. Baird, D. Baker, R. Bartlett, M. Benna, S. Bougher, D. Brain, D. Carson, S. Cauffman, P. Chamberlin, J.-Y. Chaufray, O. Cheatom, J. Clarke, J. Connerney, T. Cravens, D. Curtis, G. Delory, S. Demcak, A. DeWolfe, F. Eparvier, R. Ergun, A. Eriksson, J. Espley, X. Fang, D. Folta, J. Fox, C. Gomez-Rosa, S. Habenicht, J. Halekas, G. Holsclaw, M. Houghton,

- R. Howard, M. Jarosz, N. Jedrich, M. Johnson, W. Kasprzak, M. Kelley, T. King, M. Lankton, D. Larson, F. Leblanc, F. Lefevre, R. Lillis, P. Mahaffy, C. Mazelle, W. McClintock, J. McFadden, D. L. Mitchell, F. Montmessin, J. Morrissey, W. Peterson, W. Possel, J.-A. Sauvaud, N. Schneider, W. Sidney, S. Sparacino, A. I. F. Stewart, R. Tolson, D. Toubanc, C. Waters, T. Woods, R. Yelle, and R. Zurek. The mars atmosphere and volatile evolution (maven) mission. *Space Science Reviews*, 195(1):3–48, Dec 2015.
- [9] Mark Jesick, Stuart Demcak, Brian Young, Drew Jones, Sarah Elizabeth McCandless, and Maximilian Schadeegg. Navigation overview for the mars atmosphere and volatile evolution mission. *Journal of Spacecraft and Rockets*, 54(1):29–43, 2018/04/16 2016.
 - [10] C. G. Justus and D. L. Johnson. Mars global reference atmospheric model 2001 version (mars-gram 2001): Users guide. NASA/TM-2001-201961, NASA Marshall Spaceflight Center, April 2001.
 - [11] Richard W. Zurek, Robert H. Tolson, Darren Baird, Mark Z. Johnson, and Stephen W. Bougher. Application of maven accelerometer and attitude control data to mars atmospheric characterization. *Space Science Reviews*, 195(1):303–317, Dec 2015.
 - [12] R. W. Zurek, R. A. Tolson, S. W. Bougher, R. A. Lugo, D. T. Baird, J. M. Bell, and B. M. Jakosky. Mars thermosphere as seen in maven accelerometer data. *Journal of Geophysical Research: Space Physics*, 122(3):3798–3814, 2017. 2016JA023641.
 - [13] Scott Evans, William Taber, Theodore Drain, Jonathon Smith, Hsi-cheng Wu, Michelle Guevara, Richard Sunseri, and James Evans. MONTE: The Next Generation of Mission Design & Navigation Software. In *The 6th International Conference on Astrodynamics Tools and Techniques*, 2016.
 - [14] Drew Jones, Try Lam, Nikolas Trawny, and Clifford Lee. Using onboard telemetry for maven orbit determination. In *AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper, pages 15–202, 2015.
 - [15] Moriba K Jah, Michael E Lisano, George H Born, and Penina Axelrad. Mars aerobraking spacecraft state estimation by processing inertial measurement unit data. *Journal of guidance, control, and dynamics*, 31(6):1802–1812, 2008.